FITTING A SIMPLICIAL COMPLEX USING A VARIATION OF k-MEANS

PIOTR BEBEN

Abstract. We give a simple and effective two stage algorithm for approximating a point cloud $\mathcal{S} \subset \mathbb{R}^m$ by a simplicial complex K. The first stage is an iterative fitting procedure that generalizes k-means clustering, while the second stage involves deleting redundant simplices. A form of dimension reduction of $\mathcal{S}$ is obtained as a consequence.

arXiv:1607.03849v2 [cs.LG] 2 Aug 2016

1. Introduction

The use of simplicial complexes as a means for estimating topology has found many applications to data analysis in recent times. For example, unsupervised learning techniques such as persistent homology [4, 6] often use what are known as Cech or Vietoris-Rips filtrations to capture multi-scale topological features of a point cloud $\mathcal{S} \subset \mathbb{R}^m$. The simplicial complexes in these filtrations individually are not always a good representation of the actual physical shape of $\mathcal{S}$ since, for example, they often have a higher dimension than $\mathbb{R}^m$. Our aim is to give an algorithm that can approximate $\mathcal{S}$ to the greatest extent possible when fed any simplicial complex K mapped linearly into $\mathbb{R}^m$. This algorithm has several nice properties, including a tendency towards preserving embeddings, as well as reducing to k-means clustering when K is 0-dimensional. The resulting fitting is further refined by deleting simplices that have been poorly matched with $\mathcal{S}$; the end result being a locally linear approximation of $\mathcal{S}$. A lower dimensional representation of $\mathcal{S}$ in terms of barycentric coordinates follows by projecting onto this approximation.

2. Algorithm Description

Fix K be a (geometric) simplicial complex, and let $\{v_1, \dots, v_n\} \in K$ be the set of vertices of K. The *facets* of K are the simplices of K that have the highest dimension in the sense that they are not contained in the boundary of any other simplex. K may then be represented as a collection of facets, each of which is represented by the set of vertices that it contains. The dimension of a facet is equal to the number of vertices it contains minus 1. When we refer to the *boundary* $\partial\sigma$ of a simplex σ, we mean the union of its smaller dimensional boundary simplices, even when the simplex is embedded as a subset of $\mathbb{R}^m$. Its *interior* $int(\sigma)$ is σ minus this boundary.

Any point $x \in K$ is contained in a unique smallest dimensional simplex $\sigma_x \subseteq K$. We may represent x uniquely as a convex combination

$$x = \sum_{v_j \in \sigma_x} \lambda_{jx} v_j$$

over the vertices v_j that are contained in σ_x, where $\sum_{v_j \in \sigma_x} \lambda_{jx} = 1$ for some *barycentric coordinates* $\lambda_{jx} \geq 0$. A map $g\colon K \longrightarrow \mathbb{R}$ is said to be *linear* if it is linear on each simplex of K. Namely,

$$g(x) = \sum_{v_j \in \sigma_x} \lambda_{jx} g(v_j).$$

So g restricts to a linear embedding on each simplex of K, and is uniquely determined by the values $g(v_j)$ that it takes on its vertices v_j. Thus, we have a convenient representation of g in terms of the $g(v_j)$'s.

Key words and phrases. dimensionality reduction, k-means, clustering, unsupervised learning.

2.1. **Fitting.** Fix $\mathcal{S} \subset \mathbb{R}^m$ our finite set of data points. Suppose $f\colon K \longrightarrow \mathbb{R}^m$ is any choice of linear map, meant to represent our *initial fitting* of K to S. Starting with f, our aim is to obtain successively better fittings $f^\ell\colon K \longrightarrow \mathbb{R}^m$ of K to $\mathcal{S}$, at each iteration giving a better reflection of the shape and structure of $\mathcal{S}$. We do this as follows.

Algorithm 1 (Simplicial Means) First stage: fitting K to $\mathcal{S}$

1: set $f^0 \leftarrow f$ and $\ell \leftarrow -1$;
2: **repeat**
3: increment $\ell \leftarrow \ell + 1$;
4: **for** each $y \in \mathcal{S}$ **do**
5: find a choice of $y' \in K$ such that $f^\ell(y')$ is nearest to y, and $\sigma_y \subseteq K$ the smallest dimensional simplex that contains y';
6: compute $\lambda_{jy} \geq 0$ such that $\sum_{v_j \in \sigma_y} \lambda_{jy} = 1$ and y' is the convex combination

$$y' = \sum_{v_j \in \sigma_y} \lambda_{jy} v_j.$$

7: **end for**
8: using f^ℓ, construct a new linear map $f^{\ell+1}\colon K \longrightarrow \mathbb{R}^m$, defined on each vertex v_j by setting

$$f^{\ell+1}(v_j) \leftarrow \frac{1}{|\mathcal{N}_j^\ell|} \sum_{y \in \mathcal{N}_j^\ell} \left((1-\lambda_{jy}) f^\ell(v_j) + \lambda_{jy} y\right)$$

when $|\mathcal{N}_j^\ell| > 0$, and $f^{\ell+1}(v_j) \leftarrow f^\ell(v_j)$ otherwise, where

$$\mathcal{N}_j^\ell := \{y \in \mathcal{S} \mid v_j \in \sigma_y\}.$$

9: **until** a given stop condition has not been reached (e.g. $\max_j \|f^{\ell+1}(v_j) - f^\ell(v_j)\|$ is small)
10: set $g \leftarrow f^\ell$;
11: **return** g, each σ_y, y', and the barycentric coordinates λ_{jy}.

The $|\mathcal{N}_j^\ell| > 0$ assignment for $f^{\ell+1}(v_j)$ in Step 8 can be generalized to

$$f^{\ell+1}(v_j) \leftarrow \frac{1}{|\mathcal{N}_j^\ell|} \sum_{y \in \mathcal{N}_j^\ell} \left(\left(\frac{1-\lambda_{jy}}{1+s} \right) f^\ell(v_j) + \left(\frac{\lambda_{jy}+s}{1+s} \right) y \right)$$

where $s \geq 0$ is the *learning rate*. Larger values for s lead to faster convergence, but poorer fittings and reduced stability. Taking s to be small (≤ 0.1) but nonzero has every advantage in addition to preventing the change in the mapping of a vertex v_j becoming sluggish between iterations when the barycentric coordinates λ_{jy} are all near zero. On the other hand, this does not prevent mapping of v_j from becoming stuck in its current position when each λ_{jy} is zero (regardless of the value s) since $\mathcal{N}_j^\ell$ consists only of those $y \in \mathcal{S}$ for which y' is on the interior of some simplex σ that contains v_j (meaning those on the boundary faces of σ opposite to v_j are ignored). There are advantages and disadvantages to this. On the positive side, higher dimensional simplices are often prevented from being collapsed onto lower dimensional linear patches of data, but this can also be prevent simplices from being fitted in situations where it would be desired. For example, when fitting K homeomorphic to a sphere to data sampled from a sphere (Figure 6), *craters* can form and remain in subsequent iterations if data points surrounding the crater are in fact nearest to points lying on its rim. This issue can be circumvented (but our advantages reversed!) by taking the slightly larger set

$$\overline{\mathcal{N}}_j^\ell = \{y \in \mathcal{S} \mid y' \in \sigma \text{ for some simplex } \sigma \text{ such that } v_j \in \sigma\}$$

in place of $\mathcal{N}_j^\ell$, together with $s > 0$.

There is a straightforward intuition underlying Algorithm 1. Consider the case where at the ℓ^{th} iteration f^ℓ is a linear embedding. Then v_j and K can then be thought of as $f^\ell(v_j)$ and the subspace $f^\ell(K) \subseteq \mathbb{R}^m$, and Algorithm 1 in essence has $\mathcal{S}$ attract K towards it by having each point $y \in \mathcal{S}$ exert a pull on a nearest point $y' \in K$. The caveat here is that in doing so the embedding f^ℓ must be kept linear, so the net effect of this pull must come down to its influence on the individual vertices v_j of the simplex containing y'. The influence on each of these vertices v_j should in turn decrease with some measure of distance of y' to v_j. This is analogous to pulling on a string attached at some point along a perpendicular uniform rod floating in space, in which case the acceleration of an endpoint of the rod in the direction of the pull decreases with increasing distance from the string. Since the size or shape of the simplex in our context is irrelevant, the distance from y' to v_j is measured in terms of its barycentric coordinates λ_{jy}. In particular, if y' lies near a boundary simplex opposite to v_j, then λ_{jy} is near 0 and y has little influence on v_j, while y has full influence on v_j when $\lambda_{jy} = 1$, or equivalently, when $y' = v_j$. The accumulation of these pulling forces on a vertex v_j leads us to take the centroid of the $(1 - \lambda_{jy})v_j + \lambda_{jy}y$ over all $y \in N_j^\ell$; equivalently, over all $y \in \mathcal{S}$ that are closest to a point y' lying on the interior of a simplex that has v_j as a vertex.

When K is 0-dimensional – namely a collection of n disjoint vertices v_j – there is only a single barycentric coordinate $\lambda_{jy} = 1$ for each y' and $\mathcal{N}_j^\ell = \left\{y \in \mathcal{S} \;\middle|\; y \text{ nearest to } f^\ell(v_j)\right\}$, so Algorithm 1 reduces to the classical k-means algorithm with initial clusters $f^0(v_j)$. Algorithm 1 is thereby a high dimensional non-discrete generalization of k-means clustering. This is perhaps in the same spirit as persistent homology is a higher dimensional generalization of hierarchical clustering (also by way of simplicial complexes).

2.2. **Preserving embeddings.** To obtain a our fitting g, why not simply apply the k-means algorithm to the vertices $f(v_j)$? This would likely be a poor fitting of K since the arrangement of simplices comprising K is ignored completely. Moreover, g would probably not be an embedding irrespective of f. Take for example the case where K is the following four-vertex graph embedded in $\mathbb{R}^m = \mathbb{R}^2$

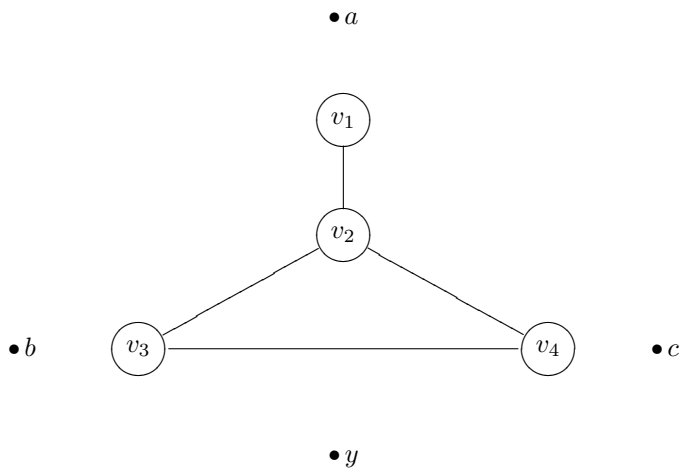

and $\mathcal{S} = \{a, b, c, y\} \subset \mathbb{R}^2$. The nearest vertex to a, b, c, and y is v_1, v_3, v_4, and v_2 respectively, so one iteration of k-means on this data results in the edge $\{v_1, v_2\}$ intersecting the edge $\{v_3, v_4\}$. On the other hand, the nearest point to y in K lies on the edge $\{v_3, v_4\}$, so a single iteration of Algorithm 1 results in K being stretched out towards $\mathcal{S}$ without violating its embedding. In fact, for non-pathological K, Algorithm 1 has a strong tendency to retain $g\colon K \longrightarrow \mathbb{R}^m$ (and each f^ℓ) as embeddings, given that the initial fitting $f\colon K \longrightarrow \mathbb{R}^m$ is an embedding.

2.3. **Refinement.** Once we have obtained a sufficiently good fitting $g\colon K \longrightarrow \mathbb{R}^m$ from Algorithm 1, the next step is to refine it by getting rid of redundant simplices. These are the simplices $\sigma \subseteq K$

whose linearly embedded image $g(\sigma)$ has no points $g(y') \in g(\sigma)$ positioned near its interior, but instead, near its boundary $\partial g(\sigma) = g(\partial\sigma)$. We then project those $g(y')$ near the boundary into one of the simplices on the boundary in order to obtain a lower dimensional approximate representation of y. This process is repeated for the projections.

Algorithm 2 Second stage: deleting redundant simplices, reducing dimension

input: a threshold value $\alpha \geq 0$, and each y', σ_y, and $g\colon K \longrightarrow \mathbb{R}^m$ from Algorithm 1;
for each $y \in \mathcal{S}$ **do**
  set $z \leftarrow y'$ and $\sigma \leftarrow \sigma_y$;
  while σ is not a vertex **do**
    find a choice of $\tilde{z}$ in $\partial\sigma$ such that $g(z)$ is nearest to $g(\tilde{z})$, and $\tilde{\sigma} \subseteq \partial\sigma$ the smallest dimensional boundary simplex that contains $\tilde{z}$;
    if $\|g(\tilde{z}) - g(z)\| \leq \alpha$ **then**
      set $z \leftarrow \tilde{z}$ and $\sigma \leftarrow \tilde{\sigma}$;
    else
      exit the while loop;
    end if
  end while
  set $\tilde{y} \leftarrow z$ and $\tilde{\sigma}_y \leftarrow \sigma$, and $\tilde{\lambda}_{jy}$ (for each $v_j \in \tilde{\sigma}_y$) the barycentric coordinates of $\tilde{y}$ in $\tilde{\sigma}_y$;
end for
set $\tilde{K} \leftarrow \bigcup_{y\in\mathcal{S}} \tilde{\sigma}_y$;
return $\tilde{K}$, and each $\tilde{y}$, $\tilde{\sigma}_y$, and $\tilde{\lambda}_{jy}$.

The output $\tilde{K}$ is a subcomplex of K with $g\colon K \longrightarrow \mathbb{R}^m$ restricting to a linear map

$$g\colon \tilde{K} \longrightarrow \mathbb{R}^m,$$

and if g is a good fitting, $g(\tilde{K})$ is a locally linear approximation of $\mathcal{S}$, and each $g(\tilde{y})$ will be a close approximation of the corresponding $y \in \mathcal{S}$. The distance threshold α is fixed for simplicity. A better choice would be to have it monotonically decrease in each iteration to counteract the effect of the distance between y' and $g(z)$ increasing in each iteration. Alternatively, a scale independent measure that is more tolerant of noise in larger scale structures might be desirable. This can be obtained for example by taking the minimum of the barycentric coordinates of y' in σ as a measure of nearness to the boundary. Algorithm 2 also always keeps those simplices that have at least one point y' near its interior with respect to α, though we could opt for something less sensitive to compromise accuracy for simplicity.

Depending on how many simplices are deleted from K to get to $\tilde{K}$, and how far down we go in projecting each $g(y')$ into boundary simplices, $\tilde{K}$ will typically have lower dimensional facets than K. At the same time, its facets and simplices typically have lower dimension than $\mathbb{R}^m$, so each y will have a lower dimensional approximate representation in terms of the barycentric coordinates $\tilde{\lambda}_{jy}$ of $\tilde{y}$ in $\tilde{\sigma}_y$. For example, if $\tilde{\sigma}_y$ happens to be a 1-dimensional simplex (edge) with vertices v_{j_1} and v_{j_2}, then $y \approx g(\tilde{y}) = tg(v_{j_1}) + (1-t)g(v_{j_2})$ where $(t, 1-t) \in \mathbb{R}^2$ are the barycentric coordinates of $\tilde{y}$ in $\tilde{\sigma}_y$. y can then be approximately represented by the barycentric coordinates $(t, 1-t)$ together with an assignment of simplex $y \mapsto \tilde{\sigma}_y$ (i.e. an assignment of those vertices in $\tilde{\sigma}_y$). In summary, points $y \in \mathcal{S}$ are assigned to linear combinations

$$\sum_{v_j \in \sigma_y} \tilde{\lambda}_{jy} g(v_j)$$

where $\tilde{\lambda}_{jy}$ are the barycentric coordinates of of $\tilde{y}$ in $\tilde{\sigma}_y$. This often gives a much more compact representation of $\mathcal{S}$, especially when m is large and large patches of points in $\mathcal{S}$ lie on approximately linear subspaces of considerably lower dimension than m.

2.4. **Some comparisons.** This approach to dimension reduction is similar to that taken by k-means and self-organizing maps (SOM) [10, 9]. The assignment of vertices (*nodes*, *neurons*, *classes*, or *clusters*, in different language) however does not have to be discrete, so the reduction is not necessarily one to dimension 0 for every $y \in \mathcal{S}$. Instead, it can be seen as a form of Fuzzy k-means [2, 20] – in this case, multiple vertices in $g(\tilde{K})$ are assigned to y to varying degrees, this depending on which simplex is nearest to y and where (in terms of barycentric coordinates) the nearest point y' lies inside this simplex. Like SOM, the mapping g is topology preserving, with adjacent vertices of K and $\tilde{K}$ tending to have nearby points in $\mathcal{S}$ assigned to them. SOM however acts directly on individual vertices by matching data points with nearest vertices instead simplices, and uses an explicit neighbourhood function on nearby vertices to effect this property. Since simplicial complexes are able to more efficiently reflect the shape of an object such as our point cloud $\mathcal{S}$, one would expect that Algorithms 1 and 2 give a better approximation of $\mathcal{S}$ than classical SOM, at the same time using fewer vertices. This is all at the expense of our dimension reduction generally being above dimension 0 depending on $\tilde{K}$.

The end result of Algorithms 1 and 2 is that local patches in $\mathcal{S}$ are approximated by the affine subspaces of $\mathbb{R}^m$ determined by the simplices of $g(\tilde{K})$. This is similar to Cluster-PCA [7], as well as the first stage of Locally Linear Embedding (LLE) [13]. In Cluster-PCA a collection of affine d-dimensional subspaces of $\mathbb{R}^m$ are iteratively fitted to their nearest points in $\mathcal{S}$ via PCA, while LLE uses convex combinations of k nearest neighbours to each point in $\mathcal{S}$ (sampled from a smoothly embedded d-manifold) in order to obtain an approximation of a local coordinate patch. In our case $\mathcal{S}$ need not be a manifold, nor does the dimension d of the affine subspaces have to be known beforehand. Instead, K has to be selected to have a high enough dimension and enough redundant facets and vertices (but not so many so as to over-fit) in order to guarantee a good approximation of $\mathcal{S}$. None-the-less, we are left with a convenient approximation of the shape of $\mathcal{S}$ in terms of $\tilde{K}$ and g. These dimension reduction techniques aside, there are various shape approximation techniques known to the computer graphics community, where, for example, a mesh is iteratively fitted to data by solving certain least-squares optimization problems [8].

2.5. **Further dimension reduction.** If $g\colon \tilde{K} \longrightarrow \mathbb{R}^m$ happens to be an embedding and $\tilde{K}$ has dimension d with d much smaller than m, then one would like to go a step further and obtain a lower dimensional embedding $h\colon \tilde{K} \longrightarrow \mathbb{R}^k$ for some $d \leq k < m$ that preserves as much of the properties of g as possible (preserving geodesic distance for instance when $\tilde{K}$ is a manifold). Assuming our fitting g is good, a dimension reduction of $\mathcal{S}$ into $\mathbb{R}^k$ follows since each $\tilde{y} \approx y \in \mathcal{S}$ and $\tilde{y}$ is in $g(\tilde{K})$.

There is a history of theoretical work dealing with the question whether *any* embedding into $\mathbb{R}^k$ of a d-dimensional simplicial complex $\tilde{K}$ can exist in the first place (ignoring g for the time-being). One of the most well-known results is Van Kampen's generalization of Kuratowski's graph planarity condition [11, 18, 14, 19], which gives an obstruction to embedding $\tilde{K}$ into $\mathbb{R}^{2d}$ lying in the degree $2d$ cohomology of a certain *deleted product* subcomplex of $\tilde{K} \times \tilde{K}$. In fact, this is the largest dimension of interest since any simplicial complex of dimension d can be embedded linearly into $\mathbb{R}^{2d+1}$ (for instance, by mapping each vertex to a unique point on the parametric curve $(t, t^2, \ldots, t^{2d+1})$). Some more recent work has focused on PL-embedding into $\mathbb{R}^{d+1}$ [3], or the tractability of embedding as a decision problem [12].

A direct computational approach is possible when $g(\tilde{K})$ is a triangulation of an embedded manifold; for example, by applying Isomap [17] to its vertices $g(v_j)$ and some choice of samples in $g(\tilde{K})$. Since the entire manifold $g(\tilde{K})$ is already given, the geodesic distances between the samples can be given exactly, and what remains is the MDS stage of Isomap. In any case, $g(\tilde{K})$ is not a sampled space, and moreover, it is often a much simpler object than $\mathcal{S}$ consisting of a much smaller collection of vertices and facets. Any dimension reduction of $\mathcal{S}$ by first reducing the dimension of the embedding of $\tilde{K}$ should be easier compared to a more head-on approach.

2.6. **Complexity.** The computationally intensive step that dominates Algorithm 1 involves finding the points $y' \in K$ for which $f^\ell(y')$ is nearest to $y \in \mathcal{S}$. These points can be found as follows.

Algorithm 3 Find nearest points on a simplex along with their barycentric coordinates

input: $\mathcal{S} \subset \mathbb{R}^m$, a simplex $\sigma \subseteq K$, and linear map $f^\ell \colon K \longrightarrow \mathbb{R}^m$;
if σ is a vertex v_j **then**
  set $\lambda^\sigma_{1y} \leftarrow 1$ and $y^\sigma \leftarrow f^\ell(v_j)$ for each $y \in \mathcal{S}$;
else
  let $r := |\mathcal{S}|$ and pick some ordering of $\mathcal{S} = \{y_1, \dots, y_r\}$;
  Let M be the $m \times d$ matrix whose i^{th} column is the vector $w_{i+1} - w_1$, where $\{w_1, \dots, w_{d+1}\} \subseteq \{f^\ell(v_1), \dots, f^\ell(v_n)\}$ are the vertices of $f^\ell(\sigma)$;
  Let S be the $m \times r$ matrix whose j^{th} column is the vector $y_j - w_1$;
  compute the Moore-Penrose pseudoinverse $\bar{M}$ of M, the $d \times r$ matrix $B := \bar{M}S$, the $1 \times r$ matrix B' whose j^{th} entry is $1 - \sum_i B_{ij}$, and take the $(d+1) \times r$ matrix $\mathcal{B} := \begin{bmatrix} B' \\ B \end{bmatrix}$;
  partition $\mathcal{S}$ into disjoint subsets $\bar{\mathcal{S}}$ and $\mathcal{S}_1, \dots, \mathcal{S}_{d+1}$ where

$$\mathcal{S}_k := \{y_j \in \mathcal{S} \mid (\mathcal{B})_{kj} < 0\} - \bigcup_{i<k} \mathcal{S}_i$$

  and

$$\bar{\mathcal{S}} := \mathcal{S} - \bigcup_i \mathcal{S}_i;$$

  for each $y_j \in \bar{\mathcal{S}}$, set $\lambda^\sigma_{iy_j} \leftarrow (\mathcal{B})_{ij}$ for $1 \le i \le d+1$, and $y^\sigma_j \leftarrow \sum_i \lambda^\sigma_{iy_j} \omega_i$;
  for each $\mathcal{S}_k$ that is non-empty **do**
    let σ_k be the boundary simplex of σ containing every vertex w_i except w_k, and set $\lambda^\sigma_{ky_j} \leftarrow 0$;
    repeat this algorithm for σ_k in place of σ and $\mathcal{S}_k$ in place of $\mathcal{S}$, letting $y^\sigma := y^{\sigma_k}$ and $\lambda^\sigma_{iy} := \lambda^{\sigma_k}_{iy}$ (for $i \neq k$) denote its output for each $y \in \mathcal{S}_k$;
  end for
end if
output: y^σ and λ^σ_{iy} for each $y \in \mathcal{S}$.

The output y^σ is the nearest point to y lying on $f^\ell(\sigma)$, and the λ^σ_{iy}'s are its barycentric coordinates in $f^\ell(\sigma)$. The smallest dimensional simplex $\sigma_y \subseteq \sigma$ for which y^σ is in $f^\ell(\sigma_y)$ can quickly be found by looking at which of the λ^σ_{iy}'s are zero.

The product $\bar{M}S$ and pseudoinverse $\bar{M}$ can be determined using singular value decomposition of M in $\mathcal{O}(dmr)$ time and $\mathcal{O}(\min\{m^2d, md^2\})$ time respectively, so all steps before the final for-loop can be executed in $\mathcal{O}(dmr + \min\{m^2d, md^2\})$. Typically $r \ge \max\{m, d\}$, in which case this simplifies to $\mathcal{O}(dmr)$. Since the recursive call in each iteration of the for-loop is executed only for those dimension $d-1$ boundary simplices $\sigma_i \subseteq \sigma$ for which $\mathcal{S}_i$ is non-empty, and since the $\mathcal{S}_i$'s are all disjoint subsets of the current input $\mathcal{S}$, then at most $r_d = \min\{r, (d+1)!\}$ recursive calls are made. So the total execution time of the above algorithm is $\mathcal{O}(dmrr_d)$.

The nearest point $f^\ell(y')$ to y is given as a choice of y^σ over all facets (not simplices) of $\sigma \subseteq K$ that minimizes $\|y^\sigma - y\|$ (y' itself is determined by the λ^σ_{iy}'s). As a result, the time complexity of each iteration of Algorithm 1 is at worst $\mathcal{O}(Ndmrr_d)$ when $r \ge \max\{m, d\}$, where N is the number of facets of K and d is the dimension of the highest dimensional facet in K. This is largely due to the cost of finding each y' and y^σ precisely. Significant improvements can be made if an approximation scheme is used. For example, since f^ℓ does not change much in consecutive iterations in Algorithm 1 for large enough ℓ, y' can be approximated in the $(\ell+1)^{th}$ iteration by searching only those facets

adjacent to the simplex that contains y' in the previous iteration. In effect, this means sending an appropriate subset $\mathcal{S}_\sigma \subseteq \mathcal{S}$ in place of $\mathcal{S}$ as input for Algorithm 3. At the same time, Algorithm 1 is easy to parallelize, since Algorithm 3 can be applied simultaneously to each facet σ of K.

3. Some Examples on Synthetic Data

Algorithm 1 was tested without the pruning stage (Algorithm 2). This was done on data $\mathcal{S}$ sampled from 1, 2, and 3 dimensional spaces using grids, triangulated meshes, or their boundaries K embedded into $\mathbb{R}^2$, $\mathbb{R}^3$, and $\mathbb{R}^4$. Precise values for nearest points were found using Algorithm 3, without any form of approximation, while the learning rate s was set at 0.1. Implementation is in MATLAB R2015b. Source code can be found at `https://github.com/pbebenSoton/smeans`.

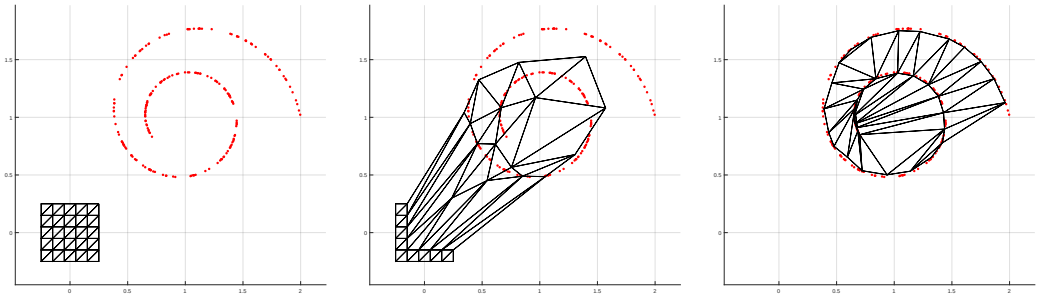

(A) 0 iterations (B) 20 iterations (C) 150 iterations

Figure 1. 200 points sampled from a *Swiss roll*, K a 1-dimensional 5×5 grid. The edges containing no points near their interior are deleted during the pruning stage.

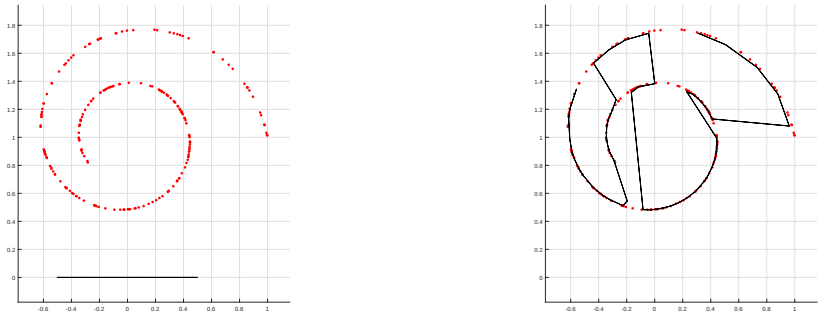

(A) 0 iterations (B) 200 iterations

Figure 2. 200 points sampled from a Swiss roll, K a line subdivided into 60 segments. The quality of this fitting is similar to classical SOM. Taking a grid of ambient dimension as in Figure 1 gives better results with fewer vertices and a similar number of facets.

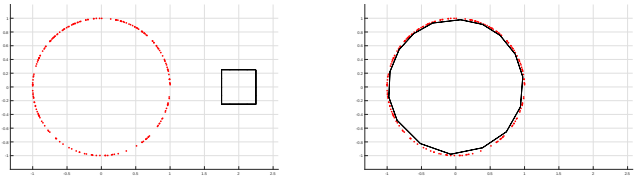

(A) 0 iterations (B) 150 iterations

FIGURE 3. 200 points sampled from a unit circle, K a square with each side divided into 4 segments. Fitted using the sets $\overline{\mathcal{N}}_j^\ell$ in place of $\mathcal{N}_j^\ell$.

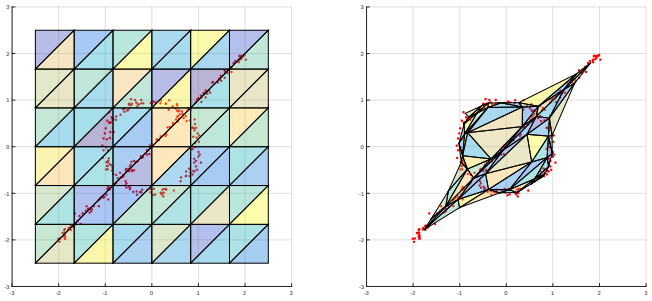

(A) 0 iterations (B) 100 iterations

FIGURE 4. 300 points sampled noisily from a circle and a line, K a triangulated 6×6 mesh.

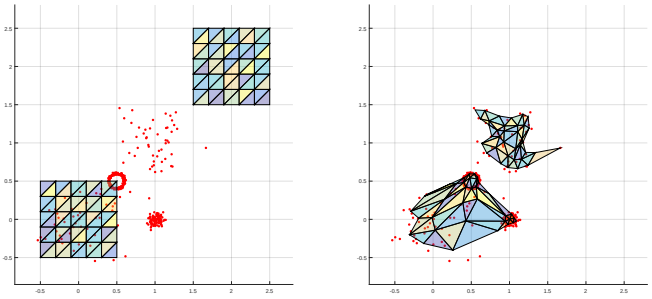

(A) 0 iterations (B) 150 iterations

FIGURE 5. A noisy sampling of 350 points containing a circle, and K a disjoint union of two 3×3 mesh *clusters*. Such clusters result in better fittings for highly disjoint data (compared to a single large mesh).

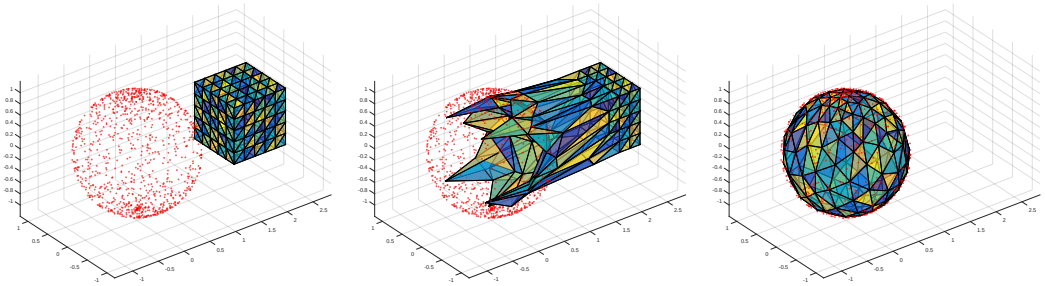

(A) 0 iterations (B) 10 iterations (C) 150 iterations

FIGURE 6. 1000 points sampled from a unit 2-sphere, K the triangulated boundary of a $5 \times 5 \times 5$ mesh. Fitted using the sets $\overline{\mathcal{N}}_j^\ell$ in place of $\mathcal{N}_j^\ell$.

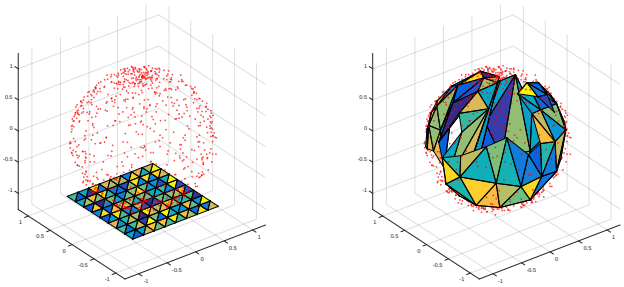

(A) 0 iterations (B) 150 iterations

FIGURE 7. 800 points noisily sampled from a unit 2-sphere, K a triangulated 8×8 mesh.

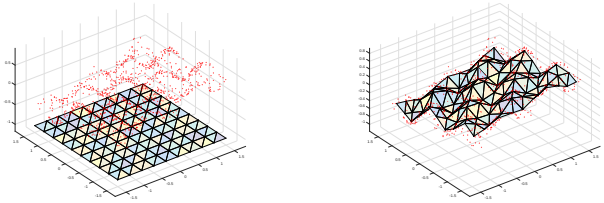

(A) 0 iterations (B) 40 iterations

FIGURE 8. 1000 points sampled from the surface $\frac{1}{3}cos(5x)sin(5y) + \frac{1}{5}(x-y)$, K a triangulated 9×9 mesh.

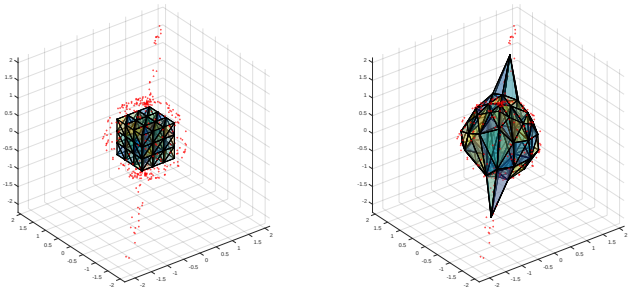

(A) 0 iterations (B) 40 iterations

FIGURE 9. 300 points noisily sampled from a 2-sphere and a line, K a triangulated $3 \times 3 \times 3$ mesh (solid cube).

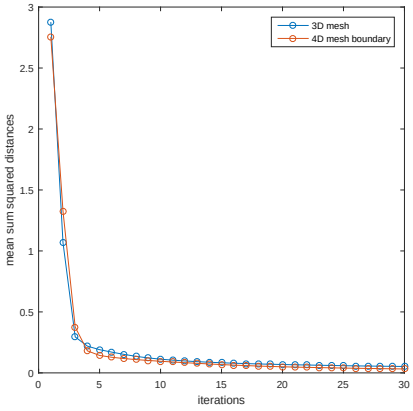

FIGURE 10. $\mathcal{S}$ taken as 2000 points noisily sampled from a unit 3-sphere embedded in 4-dimensional euclidean space $\mathbb{R}^4$, and K respectively a triangulation of a $4 \times 4 \times 4$ mesh, and the boundary of a $3 \times 3 \times 3 \times 3$ mesh. The mean sum of squared distances in each iteration is taken between points in $\mathcal{S}$ and their corresponding nearest points in the current fitting of K.

4. DISCUSSION

For non-pathological complexes K (such as above) Algorithm 1 quickly converges to a stable fitting. In most cases it also preserves the embedding of K, and where it does not, then it is usually very close to an embedding, with only a few simplices intersecting over a relatively small area (intersections are more common for less noisy data). All of these properties are less likely hold for irregularly triangulated K; for example, when the degree of some vertices is much higher than others.

If there is convergence, what is the objective function that is being minimized? Since Algorithm 1 generalizes k-means clustering, an obvious first choice is the sum of squared distances $d_\ell = \sum_{y \in \mathcal{S}} \|y - f^\ell(y')\|^2$. Indeed, there are cases where this function happens to be minimized between every successive iteration; for instance, in the example in Figures 8 and 10. On the

other hand, if we start off as in Figure 4, with $\ell = 0$ and $f^0(K)$ containing every point in $\mathcal{S}$, then $f^\ell(K)$ is unlikely to contain $\mathcal{S}$ for large enough ℓ, so d_ℓ will not always minimized. In this situation it appears more appropriate to consider the Hausdorff distance $d_H(\mathcal{S}, f^\ell(K)) = \inf\left\{\epsilon \geq 0 \;\middle|\; f^\ell(K) \subseteq (\mathcal{S})_\epsilon \text{ and } \mathcal{S} \subseteq (f^\ell(K))_\epsilon\right\}$ where A_ϵ is the ϵ-neighbourhood of a subset $A \subseteq \mathbb{R}^m$. In any case, it is not immediately clear what the objective function might be, or if there even is a single one.

Another question is: under what modifications to Algorithm 1 is preservation of embedding guaranteed for some given (or all) K? One possibility is to assume all facets of K have the same ambient dimension m as $\mathbb{R}^m$, then to update the vertices $f^\ell(v_j)$ in each iteration one at a time – at each update step considering only those $y \in \mathcal{N}_j^\ell$ that are inside a convex neighbourhood B of $f^\ell(v_j)$, where B is contained in the union of the (currently updated) facets which have $f^\ell(v_j)$ as a vertex. This can be shown to preserve embedding, it keeps vertices on the boundary fixed, and therefore points y that are inside $f^\ell(K)$ remain as such in subsequent iterations. Then in order to get a good fitting, a large enough triangulation of (say) an m-ball K is selected, together with a sufficiently expansive initial embedding f such that: there are enough interior facets in K to approximate $\mathcal{S}$. The main downside to this approach is that K would have to be infeasibly large when m is large.

As with k-means clustering, there is the issue of selecting the appropriate number of clusters (vertices) to start with. On top of this there are many other factors, such as the number, dimension, and the arrangement of facets that go into describing our simplicial complex K. Ideally, K would be taken to be a mesh of dimension d with a fine enough subdivision, embedded into $\mathbb{R}^d$ so as to contain $\mathcal{S}$. Such choice K represents a brute force approach, where the guarantee of a good fitting is the result of a high density of redundant simplices spread out uniformly in $\mathbb{R}^d$. For large d this is of course impractical – even the smallest triangulation of the 15-cube has 2^{15} vertices and over ten million facets of dimension 15, increasing exponentially with d [15]. On the other hand, a simplex of any dimension d has only one facet and $d+1$ vertices, so there is plenty of room for compromise. Skirting this issue by selecting the mesh K to have fewer subdivisions or a lower dimension can leave a large subset $\mathcal{S}'$ of points in $\mathcal{S}$ poorly approximated by the final fitting g of K (see Figures 7 and 9 for example). One way around this is to fit a new complex K' to the subset $\mathcal{S}'$, but this runs the risk of $\mathcal{S}'$ being a poor sampling of the underlying space. So there might be nothing meaningful left to fit. Instead, we can consider a *growing* procedure akin to growing SOM (GSOM) [1]: (1) first by considering the set $\mathcal{C}$ of those simplices $\sigma \subseteq K$ for which $g(\sigma)$ is a poor approximation of the points $y \in \mathcal{S}$ nearest to its interior (i.e. $\sum_{y' \in int(\sigma)} \|y - g(y')\|$ is large); (2) forming a new complex K' from K by attaching a $(k+1)$-dimensional simplex to each k dimensional simplex in $\mathcal{C}$; (3) extending g to a linear map $g' \colon K' \longrightarrow \mathbb{R}^d$; (4) then refitting K' to $\mathcal{S}$ starting with g' as our initial fitting. Simplices are the smallest objects that can be attached; a more robust alternative would be to form K' from K by attaching a triangulation of the product $L \times \Delta^1 = \left\{(x, y) \;\middle|\; x \in L, y \in \Delta^1\right\}$ to the subcomplex $L = \bigcup_{\sigma \in \mathcal{C}} \sigma$, where Δ^1 is the 1-simplex (an edge). In either approach we increase the dimension of K locally in order to capture the spread of local patches of data in $\mathcal{S}$. Then reapplying the pruning stage (Algorithm 2) just as we did before captures the underlying structure.

References

1. D. Alahakoon, S. K. Halgamuge, and B. Sirinivasan, *A self growing cluster development approach to data mining*, Proceedings of IEEE International Conference on Systems, Man and Cybernetics **3** (1998), 2901–2906.
2. James C. Bezdek, *Pattern recognition with fuzzy objective function algorithms*, Kluwer Academic Publishers, Norwell, MA, USA, 1981.
3. A. Björner and A. Goodarzi, *On Codimension one Embedding of Simplicial Complexes*, ArXiv e-prints (2016).
4. Gunnar Carlsson, *Topology and data*, Bull. Amer. Math. Soc. (N.S.) **46** (2009), no. 2, 255–308. MR 2476414
5. L. Cayton, *Algorithms for manifold learning*, Univ. of California at San Diego Tech. Rep (2005).
6. Herbert Edelsbrunner and John Harer, *Persistent homology—a survey*, Surveys on discrete and computational geometry, Contemp. Math., vol. 453, Amer. Math. Soc., Providence, RI, 2008, pp. 257–282. MR 2405684

7. K. Fukunaga and D. R. Olsen, *An algorithm for finding intrinsic dimensionality of data*, IEEE Trans. Comput. **20** (1971), no. 2, 176–183.
8. Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle, *Mesh optimization*, Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (New York, NY, USA), SIGGRAPH '93, ACM, 1993, pp. 19–26.
9. T. Kohonen, M. R. Schroeder, and T. S. Huang (eds.), *Self-organizing maps*, 3rd ed., Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2001.
10. Teuvo Kohonen, *Self-organized formation of topologically correct feature maps*, Biological Cybernetics **43** (1982), no. 1, 59–69.
11. Casimir Kuratowski, *Sur le problme des courbes gauches en topologie*, Fundamenta Mathematicae **15** (1930), no. 1, 271–283 (fre).
12. Jirí Matousek, Martin Tancer, and Uli Wagner, *Hardness of embedding simplicial complexes in* r^d, CoRR **abs/0807.0336** (2008).
13. Lawrence K. Saul and Sam T. Roweis, *Think globally, fit locally: Unsupervised learning of low dimensional manifolds*, J. Mach. Learn. Res. **4** (2003), 119–155.
14. Arnold Shapiro, *Obstructions to the imbedding of a complex in a euclidean space. I. The first obstruction*, Ann. of Math. (2) **66** (1957), 256–269. MR 0089410
15. Warren D. Smith, *A lower bound for the simplexity of then-cube via hyperbolic volumes*, European Journal of Combinatorics **21** (2000), no. 1, 131 – 137.
16. C. O. S. Sorzano, J. Vargas, and A. P. Montano, *A survey of dimensionality reduction techniques*, ArXiv e-prints (2014).
17. J. B. Tenenbaum, V. de Silva, and J. C. Langford, *A global geometric framework for nonlinear dimensionality reduction*, Science **290** (2000), 2319–2323.
18. E. R. van Kampen, *Komplexe in euklidischen Räumen*, Abh. Math. Sem. Univ. Hamburg **9** (1933), no. 1, 72–78. MR 3069580
19. Wu Wen-tsün, *A theory of imbedding, immersion, and isotopy of polytopes in a euclidean space*, Science Press, Peking, 1965. MR 0215305
20. M.-S. Yang, *A survey of fuzzy clustering*, Mathematical and Computer Modelling **18** (1993), no. 11, 1 – 16.

School of Mathematics, University of Southampton, Southampton SO17 1BJ, United Kingdom
E-mail address: P.D.Beben@soton.ac.uk; pdbcas@gmail.com